

# Bayesian estimation of non-Gaussianity in pulsar timing analysis

L. Lentati<sup>\*</sup>, M. P. Hobson, P. Alexander

*Astrophysics Group, Cavendish Laboratory, JJ Thomson Avenue, Cambridge, CB3 0HE, UK*

19 September 2014

## ABSTRACT

We introduce a method for performing a robust Bayesian analysis of non-Gaussianity present in pulsar timing data, simultaneously with the pulsar timing model, and additional stochastic parameters such as those describing red spin noise and dispersion measure variations. The parameters used to define the presence of non-Gaussianity are zero for Gaussian processes, giving a simple method of defining the strength of non-Gaussian behaviour. We use simulations to show that assuming Gaussian statistics when the noise in the data is drawn from a non-Gaussian distribution can significantly increase the uncertainties associated with the pulsar timing model parameters. We then apply the method to the publicly available 15 year Parkes Pulsar Timing Array data release 1 dataset for the binary pulsar J0437–4715. In this analysis we present a significant detection of non-Gaussianity in the uncorrelated non-thermal noise, but we find that it does not yet impact the timing model or stochastic parameter estimates significantly compared to analysis performed assuming Gaussian statistics. The methods presented are, however, shown to be of immediate practical use for current European Pulsar Timing Array (EPTA) and International Pulsar Timing Array (IPTA) datasets.

**Key words:** methods: data analysis, pulsars: general, pulsars:individual

## 1 INTRODUCTION

Millisecond pulsars (MSPs) have for some time been known to exhibit exceptional rotational stability, with decade long observations providing timing measurements with accuracies similar to atomic clocks (e.g. Kaspi, Taylor, & Ryba (1994); Matsakis, Taylor, & Ebanks (1997)). Such stability lends itself well to the pursuit of a wide range of scientific goals, e.g. observations of the pulsar PSR B1913+16 showed a loss of energy at a rate consistent with that predicted for gravitational waves (Taylor & Weisberg 1989), whilst the double pulsar system PSR J0737-3039A/B has provided precise measurements of several ‘post Keplerian’ parameters allowing for additional stringent tests of general relativity (Kramer et al. 2006).

For a detailed review of pulsar timing refer to e.g. Lorimer et al. (2004). In brief, the arrival times of pulses (TOAs) for a particular pulsar will be recorded by an observatory in a series of discrete observations over a period of time. These arrival times must all be transformed into a common frame of reference, the solar system barycenter, in order to correct for the motion of the Earth.

A model for the pulsar can then be fitted to the TOAs; this characterises the properties of the pulsar’s orbital motion, as well as its timing properties such as its orbital frequency and spin down. This is most commonly carried out using the TEMPO2 pulsar-timing packages (Hobbs, Edwards, & Manchester 2006; Edwards,

Hobbs, & Manchester 2006; Hobbs et al. 2009), or more recently, the Bayesian pulsar timing package TempoNest (Lentati et al. 2014).

When performing this fitting process, both TEMPO2 and TempoNest assume purely Gaussian statistics in the properties of the uncorrelated noise. In realistic datasets, however, this assumption is not necessarily correct. If the underlying probability density function (PDF) for the noise is not Gaussian, for example, if there is an excess of outliers relative to a Gaussian distribution, modifiers to the TOA error bars that scale their size are used to find the best approximation to a Gaussian distribution. This can be performed using a single modifier for a given receiving system determined across an entire dataset, or as in the ‘fixData’ plugin for TEMPO2 (Coles et al. 2011), where the modifier is determined separately for a series of short time lags. While the latter of these two approaches can better account for a non-Gaussian distribution in the uncorrelated noise, it does so at the expense of a potentially large number of additional free parameters, and ultimately does not address the core issue, that the underlying distribution is not Gaussian.

Both approaches then have the direct consequence of decreasing the precision with which one can estimate the timing parameters, and any other signals of interest, such as intrinsic red spin noise due to rotational irregularities in the neutron star (Shannon & Cordes 2010) or correlated noise due to a stochastic gravitational wave background (GWB) generated by, for example, coalescing black holes (e.g. Jaffe & Backer 2003; Phinney 2001). Indeed, cur-

<sup>\*</sup> E-mail: ltl21@cam.ac.uk

rently all published limits on the signals induced by a GWB have been obtained under the assumption that the statistics of the TOA errors are Gaussian (see e.g. Demorest et al. 2013; van Haasteren et al. 2011).

In this paper we introduce a method of performing a robust Bayesian analysis of non-Gaussianity present in pulsar timing data, simultaneously with the pulsar timing model, and additional stochastic parameters such as those describing the red noise, and dispersion measure variations. The parameters used to define the presence of non-Gaussianity are zero for Gaussian processes, giving a simple method of defining the strength of non-Gaussian behaviour. In Section 2 we will describe the basic principles of our Bayesian approach to data analysis, giving a brief overview of how it may be used to perform model selection, and introduce **MULTINEST**. In Sections 3 and 4 we introduce the non-Gaussian likelihood we will use in our pulsar timing analysis, and apply it to a simple toy problem. In Section 5 we then extend this likelihood to the subject of pulsar timing, and apply it to both simulated and real data in Sections 6 and 7 respectively, before finally offering some concluding remarks in Section 8.

This research is the result of the common effort to directly detect gravitational waves using pulsar timing, known as the European Pulsar Timing Array (EPTA) (Janssen et al. 2008)<sup>1</sup>.

## 2 BAYESIAN INFERENCE

Given a set of data  $D$ , Bayesian inference provides a consistent approach to the estimation of a set of parameters  $\Theta$  in a model or hypothesis  $H$ . In particular, Bayes' theorem states that:

$$\Pr(\Theta | D, H) = \frac{\Pr(D | \Theta, H)\Pr(\Theta | H)}{\Pr(D | H)}, \quad (1)$$

where  $\Pr(\Theta | D, H) \equiv \Pr(\Theta)$  is the posterior probability distribution of the parameters,  $\Pr(D | \Theta, H) \equiv L(\Theta)$  is the likelihood,  $\Pr(\Theta | H) \equiv \pi(\Theta)$  is the prior probability distribution, and  $\Pr(D | H) \equiv Z$  is the Bayesian Evidence.

Since the evidence is independent of the parameters  $\Theta$  it is typically ignored when one is only interested in performing parameter estimation. In this case inferences are obtained by taking samples from the (unnormalised) posterior using, for example, standard Markov chain Monte Carlo (MCMC) sampling methods.

For model selection, however, the evidence is key, and is defined simply as the factor required to normalise the posterior over  $\Theta$ :

$$Z = \int L(\Theta)\pi(\Theta)d^n\Theta, \quad (2)$$

where  $n$  is the dimensionality of the parameter space.

As the evidence is just the average of the likelihood over the prior, it will be larger for a simpler model with a compact parameter space if more of that parameter space is likely. More complex models where large areas of parameter space have low likelihood values will have a smaller evidence even if the likelihood function is very highly peaked, unless they are significantly better at explaining the data. Thus, the evidence automatically implements Occam's razor.

The question of model selection between two models  $H_0$  and  $H_1$  can be answered via the model selection ratio  $R$ , commonly referred to as the 'Bayes Factor':

$$R = \frac{\Pr(H_1 | D)}{\Pr(H_0 | D)} = \frac{\Pr(D | H_1)\Pr(H_1)}{\Pr(D | H_0)\Pr(H_0)} = \frac{Z_1 \Pr(H_1)}{Z_0 \Pr(H_0)}, \quad (3)$$

where  $\Pr(H_1)/\Pr(H_0)$  is the a priori probability ratio for the two models, which in this work we will set to unity but occasionally requires further consideration.

The Bayes factor then allows us to obtain the probability of one model compared the other simply as:

$$P = \frac{R}{1 + R}. \quad (4)$$

In practice when performing Bayesian analysis we do not work with the likelihood, but the log likelihood. In this case the quantity of interest is the log Bayes Factor, which is simply the difference in the log evidence for the two models. For example, a difference in the log evidence of 3 for two competing models gives a Bayes factor of  $\sim 20$ , which in turn gives a probability of  $\sim 95\%$ . We use the difference in the log evidence in Sections 6 and 7 to perform model selection between our Gaussian and non-Gaussian models.

### 2.1 Nested sampling and evidence evaluation

While many techniques exist for calculating the evidence, such as thermodynamic integration (O'Ruanidh & Fitzgerald 1996), it remains a challenging task both numerically and computationally, with evidence evaluation at least an order-of-magnitude more costly than parameter estimation.

Nested sampling (Skilling 2004) is an approach designed to make the calculation of the evidence more efficient, and also produces posterior inferences as a by-product. The **MULTINEST** algorithm (Feroz, Hobson, & Bridges 2009; Feroz & Hobson 2008) builds upon this nested sampling framework, and provides an efficient means of sampling from posteriors that may contain multiple modes and/or large (curving) degeneracies, and also calculates the evidence. Since its release **MULTINEST** has been used successfully in a wide range of astrophysical problems, including inferring the properties of a potential stochastic gravitational wave background in pulsar timing array data (Lentati et al. 2013), and is also used in the Bayesian pulsar timing package **TEMPO****NEST**. This technique has greatly reduced the computational cost of Bayesian parameter estimation and model selection, and is employed in this paper.

## 3 A NON-GAUSSIAN LIKELIHOOD

In this section we will outline the method adopted for including non-Gaussian behaviour in our analysis. We use the approach developed in Rocha et al. (2001), which is based on the energy eigenmode wavefunctions of a simple harmonic oscillator. We will describe this in brief below in order to aid future discussion.

We begin by considering our data, the vector  $\mathbf{d}$  of length  $N_d$ , as the sum of some signal  $\mathbf{s}$  and noise  $\mathbf{n}$  such that:

$$\mathbf{d} = \mathbf{s} + \mathbf{n}. \quad (5)$$

We can then construct the likelihood that the residuals after subtracting our model signal from the data follows an uncorrelated Gaussian distribution of width  $\sigma$  as:

$$\Pr(\mathbf{d}|\sigma) = \frac{1}{\sqrt{(2\pi)^{N_d}\det(\mathbf{N})}} \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{s})^T \mathbf{N}^{-1}(\mathbf{d} - \mathbf{s})\right], \quad (6)$$

with  $\mathbf{N}$  the diagonal noise covariance matrix for the residuals, such that  $N_{ii} = \sigma^2$ , and  $\det(\mathbf{N})$  the determinant of  $\mathbf{N}$ .

<sup>1</sup> [www.epta.eu.org/](http://www.epta.eu.org/)

We now extend this to the general case in order to allow for non-Gaussian distributions by modelling our PDF as the sum of a set of Gaussians, modified by Hermite polynomials  $H_n(x)$  (see e.g. Hall (1989) for previous uses of Hermite polynomials in describing departures from Gaussianity), defined as:

$$H_n(x) = (-1)^n \exp(x^2) \frac{d^n}{dx^n} \exp(-x^2). \quad (7)$$

Therefore, for a general random variable  $x$  the PDF for fluctuations in  $x$  can be written:

$$\Pr(x|\sigma, \alpha) = \exp\left[-\frac{x^2}{2\sigma^2}\right] \left| \sum_{n=0}^{\infty} \alpha_n C_n H_n\left(\frac{x}{\sqrt{2}\sigma}\right) \right|^2 \quad (8)$$

with  $\alpha_n$  free parameters that describe the relative contributions of each term to the sum, and

$$C_n = \frac{1}{(2^n n! \sqrt{2\pi}\sigma)^{1/2}}, \quad (9)$$

is a normalization factor. Equation 8 forms a complete set of PDFs, normalised such that:

$$\int_{-\infty}^{\infty} dx \exp\left[-\frac{x^2}{\sigma^2}\right] C_n H_n\left(\frac{x}{\sqrt{2}\sigma}\right) C_m H_m\left(\frac{x}{\sqrt{2}\sigma}\right) = \delta_{nm}, \quad (10)$$

with  $\delta_{nm}$  the Kronecker delta, where the ground state,  $H_0$ , reproduces a standard Gaussian PDF, and any non-Gaussianity in the distribution of  $x$  will be reflected in non-zero values for the coefficients  $\alpha_n$  associated with higher order states.

The only constraint we must place on the values of the amplitudes  $\alpha$  is:

$$\sum_{n=0}^{n_{\max}} |\alpha_n|^2 = 1 \quad (11)$$

with  $n_{\max}$  the maximum number of coefficients to be included in the model for the PDF. This is performed most simply by setting:

$$\alpha_0 = \sqrt{1 - \sum_{n=1}^{n_{\max}} |\alpha_n|^2}. \quad (12)$$

We can therefore rewrite Eq. 6 in this more general form as:

$$\begin{aligned} \Pr(\mathbf{d}|\sigma, \alpha) &= \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{s})^T \mathbf{N}^{-1}(\mathbf{d} - \mathbf{s})\right] \\ &\times \prod_{i=1}^{N_d} \left| \sum_{n=0}^{n_{\max}} \alpha_n C_n H_n\left(\frac{d_i - s_i}{\sqrt{2}\sigma}\right) \right|^2. \end{aligned} \quad (13)$$

The advantage of this method is that one may use a finite set of non-zero  $\alpha_n$  to model the non-Gaussianity, without mathematical inconsistency. Any truncation of the series still yields a proper distribution, in contrast to the more commonly used Edgeworth expansion (e.g. Contaldi et al. 2000).

#### 4 APPLICATION TO A TOY PROBLEM

Before applying the formalism described in Section 3 to the practice of pulsar timing, we first demonstrate its use in a toy problem. Here our data vector  $\mathbf{d}$  contains 10000 points drawn from a non-Gaussian distribution obtained using Eq. 8, with parameters listed in Table 1.

We then sample over the 4 dimensional problem using MULTINEST using Eq. 13 as our likelihood, with  $\mathbf{s} = 0$ . We note here that in practice when dealing with real data the correct number of coefficients to use in the analysis will not be a known quantity.

**Table 1.** Parameters used to generate non-Gaussian noise in a simple toy problem.

Parameter	Value	Parameter Estimate
$\sigma$	1	$0.997 \pm 0.005$
$\alpha_1$	0.1	$0.105 \pm 0.006$
$\alpha_2$	0.2	$0.198 \pm 0.006$
$\alpha_3$	0.4	$0.402 \pm 0.005$

In this case one could take the approach of performing the analysis multiple times including an increasing number of non-Gaussian terms in order to find the value that maximises the evidence, an approach we take in Section 7.

Fig. 1 (left) shows the distribution of points in the vector  $\mathbf{d}$  (red histogram), along with the mean model solution obtained for the distribution in our analysis (smooth blue line) which we also give in Table 1 along with the value for one standard deviation in each parameter. Fig. 1 (right) shows the one and two-dimensional marginalised posterior distributions for the parameters sampled in the toy problem, with vertical lines indicating the values used in the simulation. We find that the analysis successfully reproduces the PDF for the values in  $\mathbf{d}$ .

#### 5 NON-GAUSSIAN PULSAR TIMING LIKELIHOOD

For any pulsar we can, as in Section 3, write the TOAs for the pulses as a sum of both a deterministic and a stochastic component:

$$\mathbf{d} = \mathbf{s} + \mathbf{n}, \quad (14)$$

where  $\mathbf{d}$  represents the  $N_d$  TOAs for a single pulsar, with  $\mathbf{s}$  and  $\mathbf{n}$  the deterministic and stochastic contributions to the total respectively.

Writing the deterministic signal due to the timing model parameters  $\epsilon$  as  $\tau(\epsilon)$ , and the uncertainty associated with a particular TOA  $i$  as  $\sigma_i$  we can rewrite our non-Gaussian likelihood from Eq. 13 simply as:

$$\begin{aligned} \Pr(\mathbf{d}|\epsilon, \beta, \alpha) &= \exp\left[-\frac{1}{2}(\mathbf{d} - \tau(\epsilon))^T \mathbf{N}^{-1}(\mathbf{d} - \tau(\epsilon))\right] \\ &\times \prod_{i=1}^{N_d} \left| \sum_{n=0}^{n_{\max}} \alpha_n C_{i,n} H_n\left(\frac{d_i - \tau(\epsilon)_i}{\sqrt{2}\beta\sigma_i}\right) \right|^2, \end{aligned} \quad (15)$$

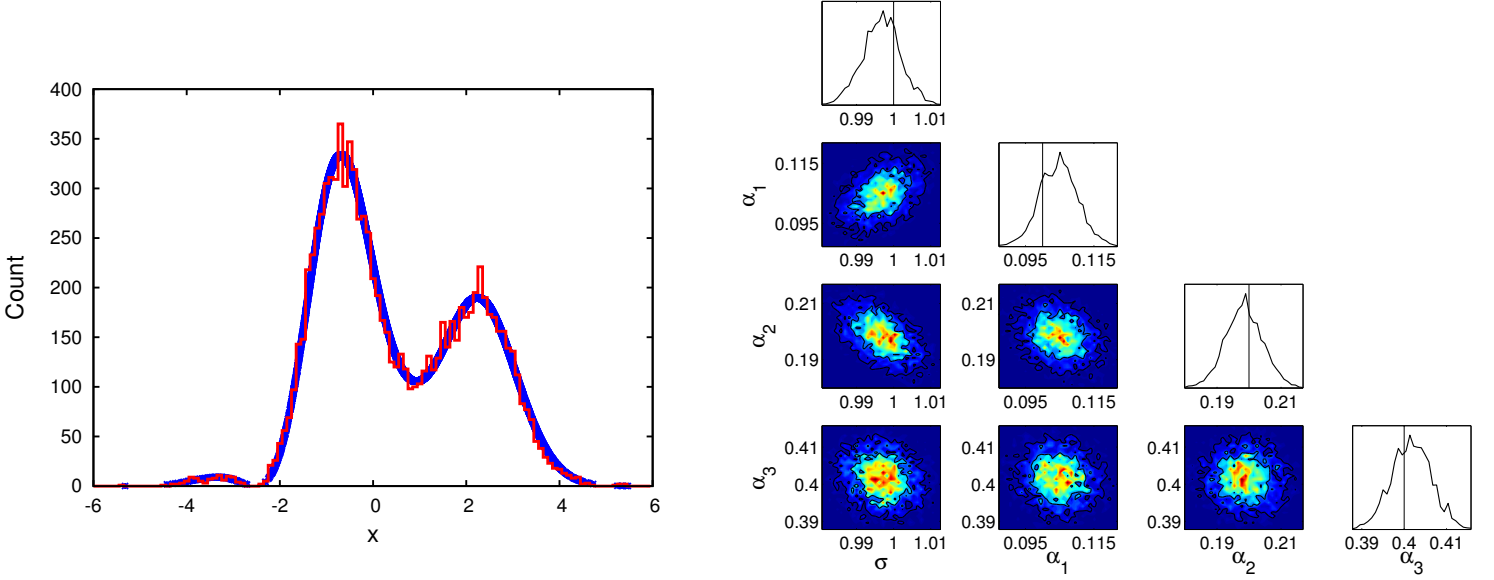
where we have redefined the constant  $C$  as:

$$C_{i,n} = \frac{1}{(2^n n! \sqrt{2\pi}\beta\sigma_i)^{1/2}}, \quad (16)$$

to allow for the more general case where each TOA  $d_i$  has a different uncertainty  $\sigma_i$ , and we include the scaling factor  $\beta$ , commonly referred to as ‘EFAC’, as a free parameter to modify these values to account for possible mis-calibration of this radiometer noise.

##### 5.1 Additional white noise

In typical pulsar timing analysis the white noise is considered the sum of multiple terms, the radiometer noise associated with a given TOA discussed in the previous section, and additional terms that add in quadrature and represent sources of time independent noise. These can include, for example, contributions from the high frequency tail of the pulsar’s red spin noise power spectrum, or, jitter noise that results from the time averaging of a finite number of



**Figure 1.** (left) The distribution of points in the vector  $\mathbf{d}$  (red histogram), and the mean model solution obtained for the distribution in our analysis (smooth blue line). (right) The one and two-dimensional marginalised posterior distributions for the parameters sampled in the toy problem; vertical lines represent the values used in the simulations. We find that the analysis successfully reproduces the target distribution.

single pulses to form each TOA (see e.g. Shannon et al. (2014)). For simplicity we will refer to these quadrature terms simply as ‘EQUAD’.

For any given TOA the total white noise can therefore be considered to be the sum of these separate processes. In incorporating the different white noise processes into our likelihood we first note that, for a random variable  $c$ , defined as  $c = a + b$  with  $a$  and  $b$  themselves random variables, we can write the PDF for  $c$  as:

$$\Pr(c) = \Pr(a) * \Pr(b), \quad (17)$$

where  $\Pr(a)$  and  $\Pr(b)$  are the PDFs for variables  $a$  and  $b$  respectively, and  $*$  represents a convolution. Using this result we can then take one of two approaches. Firstly, we can use Eq. 15 and simply fit for a single non-Gaussian probability density function for the combined radiometer and quadrature terms by taking the total uncertainty to be:

$$\hat{\sigma}_i^2 = \beta^2 \sigma_i^2 + \gamma^2 \quad (18)$$

where  $\gamma$  represents the quadrature component of the total error bar in which case we simply replace  $\beta\sigma_i$  with  $\hat{\sigma}_i$  in Eq. 15.

Alternatively we can allow the PDFs for the radiometer and quadrature terms to be different, for example, by assuming the radiometer term is Gaussian, and that any non-Gaussianity comes from all, or a subset of, the terms added in quadrature. In this case we will have:

$$\Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta, \gamma, \boldsymbol{\alpha}) = \Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta) * \Pr(\mathbf{d}|\boldsymbol{\epsilon}, \gamma, \boldsymbol{\alpha}) \quad (19)$$

where

$$\Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta) = \frac{1}{\sqrt{(2\pi)^{N_d} |\mathbf{N}_\beta|}} \exp \left[ -\frac{1}{2} (\mathbf{d} - \tau(\boldsymbol{\epsilon}))^T \mathbf{N}_\beta^{-1} (\mathbf{d} - \tau(\boldsymbol{\epsilon})) \right], \quad (20)$$

with  $\mathbf{N}_{\beta,i,i} = \beta^2 \sigma_i^2$ , and

$$\begin{aligned} \Pr(\mathbf{d}|\boldsymbol{\epsilon}, \gamma, \boldsymbol{\alpha}) &= \exp \left[ -\frac{1}{2} (\mathbf{d} - \tau(\boldsymbol{\epsilon}))^T \mathbf{J}^{-1} (\mathbf{d} - \tau(\boldsymbol{\epsilon})) \right] \\ &\times \prod_{i=1}^{N_d} \left[ \sum_{n=0}^{n_{\max}} \alpha_n C_{i,n,\gamma} H_n \left( \frac{d_i - \tau(\boldsymbol{\epsilon})_i}{\sqrt{2}\gamma} \right) \right]^2. \end{aligned}$$

with  $\mathbf{J}_{i,i} = \gamma^2$ .

## 5.2 Sampling from the un-marginalised EQUAD

A second approach to including both EFAC and EQUAD parameters in the likelihood is to parameterise the EQUAD term using  $N_d$  free parameters  $j_i$  that each represent a shift in a given TOA  $i$ , and then having a prior on those parameters that describes their underlying distribution.

If we first assume a Gaussian distribution on the parameters  $\mathbf{j}$  we can write our likelihood as:

$$\Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta, \mathbf{j}, \mathbf{J}) = \Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta, \mathbf{j}) \times \Pr(\mathbf{j}|\mathbf{J}) \quad (21)$$

with,

$$\begin{aligned} \Pr(\mathbf{d}|\boldsymbol{\epsilon}, \beta, \mathbf{j}) &= \frac{1}{\sqrt{(2\pi)^{N_d} \det \mathbf{N}_\beta}} \\ &\times \exp \left[ -\frac{1}{2} (\mathbf{d} - \tau(\boldsymbol{\epsilon}) - \mathbf{j})^T \mathbf{N}_\beta^{-1} (\mathbf{d} - \tau(\boldsymbol{\epsilon}) - \mathbf{j}) \right], \end{aligned} \quad (22)$$

and

$$\Pr(\mathbf{j}|\mathbf{J}) = \frac{1}{\sqrt{(2\pi)^{N_d} \det \mathbf{J}}} \exp \left[ -\frac{1}{2} \mathbf{j}^T \mathbf{J}^{-1} \mathbf{j} \right], \quad (23)$$

with  $\mathbf{J}$  a diagonal matrix, where each element is equal to  $\gamma^2$  as defined previously.

This formalism is infact completely equivalent to the convolution described in Section 5.1, which can be seen quite transparently by simply considering the definition of the convolution for two functions  $P$  and  $Q$ :

$$\int_{-\infty}^{\infty} P(\mathbf{d} - \mathbf{j})Q(\mathbf{j})d\mathbf{j} = P(\mathbf{d}) * Q(\mathbf{d}). \quad (24)$$

By equating  $Q(\mathbf{j})$  to our prior on the parameters  $\mathbf{j}$ , and  $P(\mathbf{d} - \mathbf{j})$  to our likelihood, the integration over all  $\mathbf{j}$ 's is seen to be equivalent to a convolution.

For completeness, we will now show that by re-marginalising over the parameters  $\mathbf{j}$  we recover the definition of the white noise given in Eqn 18 in the Gaussian case. In order to perform this marginalisation, we first write the log of the likelihood in Eq 21, which denoting  $(\mathbf{N}_\beta^{-1} + \mathbf{J}^{-1})$  as  $\mathbf{\Sigma}$  and  $\mathbf{N}_\beta^{-1}\mathbf{d}$  as  $\bar{\mathbf{d}}$  is given by:

$$\log L = -\frac{1}{2}\mathbf{d}^T\bar{\mathbf{d}} - \frac{1}{2}\mathbf{j}^T\mathbf{J}^{-1}\mathbf{j} + \bar{\mathbf{d}}^T\mathbf{j}. \quad (25)$$

Taking the derivative of  $\log L$  with respect to  $\mathbf{j}$  gives us:

$$\frac{\partial \log L}{\partial \mathbf{j}} = -\mathbf{\Sigma}\mathbf{j} + \bar{\mathbf{d}}, \quad (26)$$

which can be solved to give us the maximum likelihood vector of coefficients  $\hat{\mathbf{j}}$ :

$$\hat{\mathbf{j}} = \mathbf{\Sigma}^{-1}\bar{\mathbf{d}}. \quad (27)$$

Re-expressing Eq. 25 in terms of  $\hat{\mathbf{j}}$ :

$$\log L = -\frac{1}{2}\mathbf{d}^T\bar{\mathbf{d}} + \frac{1}{2}\hat{\mathbf{j}}^T\mathbf{\Sigma}\hat{\mathbf{j}} - \frac{1}{2}(\mathbf{j} - \hat{\mathbf{j}})^T\mathbf{\Sigma}(\mathbf{j} - \hat{\mathbf{j}}), \quad (28)$$

the 3rd term in this expression can then be integrated with respect to the  $N_d$  elements in  $\mathbf{j}$  to give:

$$\begin{aligned} I &= \int_{-\infty}^{+\infty} d\mathbf{j} \exp\left[-\frac{1}{2}(\mathbf{j} - \hat{\mathbf{j}})^T\mathbf{\Sigma}(\mathbf{j} - \hat{\mathbf{j}})\right] \\ &= (2\pi)^{N_d} \det \mathbf{\Sigma}^{-\frac{1}{2}}. \end{aligned} \quad (29)$$

Our marginalised likelihood for EQUAD is then given as:

$$\log L = -\frac{1}{2}|\mathbf{\Sigma}| - \frac{1}{2}|\mathbf{J}| - \frac{1}{2}|\mathbf{N}_\beta| - \frac{1}{2}(\mathbf{d}^T\bar{\mathbf{d}} - \bar{\mathbf{d}}^T\mathbf{\Sigma}^{-1}\bar{\mathbf{d}}). \quad (30)$$

Given the definitions of  $\bar{\mathbf{d}}$ ,  $\mathbf{J}$ , and  $\mathbf{\Sigma}$  we can write for a single TOA  $i$ :

$$\begin{aligned} (\mathbf{d}^T\bar{\mathbf{d}} - \bar{\mathbf{d}}^T\mathbf{\Sigma}^{-1}\bar{\mathbf{d}})_i &= d_i \left( \frac{1}{\beta^2\sigma_i^2} - \frac{1}{\beta^2\sigma_i^2} \left( \frac{1}{\beta^2\sigma_i^2} + \frac{1}{\gamma^2} \right)^{-1} \frac{1}{\beta^2\sigma_i^2} \right) d_i, \\ &= \frac{d_i^2}{\beta^2\sigma_i^2 + \gamma^2} \end{aligned} \quad (31)$$

and similarly for the determinants.

We can use this un-marginalised distribution to trivially include non-Gaussianity in the EQUAD term only by altering the prior term:

$$\Pr(\mathbf{j}|\mathbf{J}, \alpha) = \exp\left[-\frac{1}{2}\mathbf{j}^T\mathbf{J}^{-1}\mathbf{j}\right] \prod_{i=1}^{N_d} \left[ \sum_{n=0}^{n_{\max}} \alpha_n C_{i,n,\gamma} H_n\left(\frac{j_i}{\sqrt{2\gamma}}\right) \right]^2. \quad (32)$$

It is still possible to then analytically marginalise over the  $j$  parameters, which will result in the convolution of a Gaussian and non-Gaussian probability density function.

### 5.3 Analytic marginalisation over the EFAC parameters

An interesting possibility that arises when parameterising the EQUAD parameter as in Section 5.2 is that it allows us to marginalise analytically over the EFAC white noise parameters. If we assume a uniform prior on the amplitude of the EFAC parameter  $\beta$  and define:

$$r^2 = \sum_{i=1}^{N_d} (d_i - \tau(\epsilon)_i - j_i)^2 / \sigma_i^2 \quad (33)$$

we can write our probability distribution as:

$$\Pr(\mathbf{d}|\epsilon, \mathbf{j}, \beta) \propto \frac{1}{\beta^{N_d}} \exp\left[-\frac{r^2}{2\beta^2}\right] \quad (34)$$

which can then be integrated over  $\beta$  between some finite lower limit  $a$  and a finite upper limit  $b$ . While this does have an analytic solution given by:

$$\int_a^b d\beta \frac{1}{\beta^{N_d}} \exp\left[-\frac{r^2}{2\beta^2}\right] \propto \quad (35)$$

$$\left(\frac{1}{r^2}\right)^{\frac{N_d-1}{2}} \left( \Gamma\left[\frac{N_d-1}{2}, \frac{r^2}{2b^2}\right] - \Gamma\left[\frac{N_d-1}{2}, \frac{r^2}{2a^2}\right] \right) \quad (36)$$

with  $\Gamma$  the upper incomplete Gamma function, if the lower bound for the integral  $a$  is non zero, the precision required to carry out the difference can be substantial for large  $N_d$ . In this case it may be preferable simply to integrate using a standard numerical integration package, where the integral is rephrased to make it computable at double precision as:

$$\Pr(\mathbf{d}|\epsilon, \mathbf{j}, \beta) \propto \frac{\exp(M)}{\beta^{N_d}} \exp\left[-\frac{r^2}{2\beta^2}\right], \quad (37)$$

with  $M$  the maximum value taken by the log of Eq. 34.

If one takes the case where the lower bound on the integral is 0, the integral in Eq 35 simplifies to:

$$\int_0^b d\beta \frac{1}{\beta^{N_d}} \exp\left[-\frac{r^2}{2\beta^2}\right] \propto \left(\frac{1}{r^2}\right)^{\frac{N_d-1}{2}} \Gamma\left[\frac{N_d-1}{2}, \frac{r^2}{2b^2}\right], \quad (38)$$

however, if the upper limit is taken to be arbitrarily large, but still finite, the log of the Gamma function term is effectively constant, and the result is a trivial, scale invariant log likelihood:

$$\log L = -\frac{N_d-1}{2} \log r^2. \quad (39)$$

We note that in principle if a Gaussian prior is assumed on the EQUAD parameters  $\mathbf{j}$  one could also marginalise analytically over the hyper parameter  $\gamma$ . In this case we would redefine:

$$r^2 = \sum_{i=1}^{N_d} (j_i)^2 \quad (40)$$

giving us a probability distribution:

$$\Pr(\mathbf{j}|\mathbf{J}) \propto \frac{1}{\gamma^{N_d}} \exp\left[-\frac{r^2}{2\gamma^2}\right] \quad (41)$$

from where we can proceed as before. In practice however we find this approach less efficient to sample from in comparison to the convolved likelihood described in Section 5.1, and so we take the latter approach in Sections 6 and 7.

### 5.4 Additional red spin noise

In order to include additional red noise processes we begin by taking the same approach as that given in Lentati et al. (2014), which we will describe in brief below to aid further discussion. Writing the red noise component of the stochastic signal, which we will denote  $\mathbf{d}_{\text{red}}$ , in terms of its Fourier coefficients  $\mathbf{a}_{\text{red}}$  so that  $\mathbf{d}_{\text{red}} = \mathbf{F}_{\text{red}} \mathbf{a}_{\text{red}}$  where  $\mathbf{F}_{\text{red}}$  denotes the Fourier transform such that for frequency  $\nu$  and time  $t$  we will have both:

$$F_{\text{red}}(\nu, t) = \frac{1}{T} \sin(2\pi\nu t), \quad (42)$$

and an equivalent cosine term. Here  $T$  represents the total observing span for the pulsar, and  $\nu$  the frequency of the signal to be sampled. Defining the number of coefficients to be sampled by  $n_{\text{red}}$ , we can then include the set of frequencies with values  $n/T$ , where  $n$  extends from 1 to  $n_{\text{red}}$ . For typical PTA data Lee et al. (2012) show that a low frequency cut off of  $1/T$  is sufficient to accurately describe the expected long term variations present in the data. If necessary though it is also possible to specify arbitrary sets of frequencies such that terms with  $\nu \ll 1/T$  can be included in the model, or to allow noise terms where the frequency itself is a free parameter.

For a single pulsar the covariance matrix  $\boldsymbol{\varphi}_{\text{red}}$  of the Fourier coefficients  $\mathbf{a}_{\text{red}}$  will be diagonal, with components

$$\varphi_{\text{red},ij} = \langle a_{\text{red},i} a_{\text{red},j}^* \rangle = \varphi_{\text{red},i} \delta_{ij}, \quad (43)$$

where there is no sum over  $i$ , and the brackets  $\langle \dots \rangle$  denotes the expectation value such that the set of coefficients  $\{\varphi_i\}$  represent the theoretical power spectrum of the red noise signal present in the timing data.

Whilst Eq. 43 states that the Fourier modes are orthogonal to one another, this does not mean that we assume they are orthogonal in the time domain where they are sampled, and it can be shown that this non-orthogonality is accounted for within the likelihood. Instead, in Bayesian terms, Eq. 43 represents our prior knowledge of the power spectrum coefficients within the data. We are therefore stating that, whilst we do not know the form the power spectrum will take, we know that the underlying Fourier modes are still orthogonal by definition, regardless of how they are sampled in the time domain. It is here then that, should one wish to fit a specific model to the power spectrum coefficients at the point of sampling, such as a broken, or single power law, the set of coefficients  $\{\varphi_i\}$  should be given by some function  $f(\Theta)$ , where we sample from the parameters  $\Theta$  from which the power spectrum coefficients  $\{\varphi_i\}$  can then be derived.

We can then use the signal realisation of the red noise process given by  $\mathbf{F}\mathbf{a}$  to alter the model TOAs given by the timing model,  $\tau(\epsilon)$  such that:

$$\hat{\tau}(\epsilon, \mathbf{a}_{\text{red}}) = \tau(\epsilon) - \mathbf{F}_{\text{red}} \mathbf{a}_{\text{red}}, \quad (44)$$

enabling us to write the joint probability density  $\Pr(\epsilon, \beta, \alpha, \boldsymbol{\varphi}, \mathbf{a}_{\text{red}} | \mathbf{d})$ , as:

$$\begin{aligned} \Pr(\epsilon, \beta, \alpha, \boldsymbol{\varphi}, \mathbf{a}_{\text{red}} | \mathbf{d}) &\propto \Pr(\mathbf{d} | \epsilon, \beta, \alpha, \mathbf{a}_{\text{red}}) \\ &\times \Pr(\mathbf{a}_{\text{red}} | \boldsymbol{\varphi}_{\text{red}}) \Pr(\boldsymbol{\varphi}_{\text{red}}). \end{aligned} \quad (45)$$

For our choice of  $\Pr(\boldsymbol{\varphi}_{\text{red}})$  we use an uninformative prior that is uniform in  $\log_{10}$  space, and draw our samples from the parameter  $\rho_{\text{red},i} = \log_{10}(\varphi_{\text{red},i})$  instead of  $\varphi_{\text{red},i}$ . Given this choice of prior the conditional distributions that make up Eq. 45 can be written:

$$\begin{aligned} \Pr(\mathbf{d} | \epsilon, \beta, \alpha, \mathbf{a}_{\text{red}}) &= \exp \left[ -\frac{1}{2} (\mathbf{d} - \hat{\tau}(\epsilon))^T \mathbf{N}^{-1} (\mathbf{d} - \hat{\tau}(\epsilon)) \right] \\ &\times \prod_{i=1}^{N_d} \left| \sum_{n=0}^{n_{\text{max}}} \alpha_n C_{i,n} H_n \left( \frac{d_i - \hat{\tau}(\epsilon)_i}{\sqrt{2}\beta\sigma_i} \right) \right|^2, \end{aligned} \quad (46)$$

and:

$$\Pr(\mathbf{a}_{\text{red}} | \boldsymbol{\rho}_{\text{red}}) \propto \frac{1}{\sqrt{\det \boldsymbol{\varphi}_{\text{red}}}} \exp \left[ -\frac{1}{2} \mathbf{a}_{\text{red}}^T \boldsymbol{\varphi}_{\text{red}}^{-1} \mathbf{a}_{\text{red}} \right]. \quad (47)$$

As in section 5.2 we note that we need not use a Gaussian prior on the Fourier coefficients. If we assume that the red noise follows a power law, such that the matrix  $\boldsymbol{\varphi}$  is a function of an amplitude  $A_{\text{red}}$  and spectral index  $\kappa_{\text{red}}$ :

$$\varphi_{ii} = \frac{A_{\text{red}}^2}{12\pi^2 T_s} \left( \frac{1}{1\text{yr}} \right)^{-3} (1\text{yr } \nu_i)^{-\kappa_{\text{red}}}, \quad (48)$$

with  $T_s$  the total observing span of the dataset in seconds, and  $\nu_i$  the frequency of the power spectrum coefficient  $\varphi_{ii}$ , then we can parameterise the non-Gaussianity in the coefficients  $\mathbf{a}$  as before:

$$\begin{aligned} \Pr(\mathbf{a}_{\text{red}} | \boldsymbol{\rho}_{\text{red}}, \alpha) &= \exp \left[ -\frac{1}{2} \mathbf{a}_{\text{red}}^T \boldsymbol{\varphi}_{\text{red}}^{-1} \mathbf{a}_{\text{red}} \right] \\ &\times \prod_{i=1}^{n_{\text{red}}} \left| \sum_{n=0}^{n_{\text{max}}} \alpha_n C_{i,n,\varphi_{\text{red},i}} H_n \left( \frac{a_i}{\sqrt{2\varphi_{\text{red},i}}} \right) \right|^2. \end{aligned} \quad (49)$$

We note here that, if a Gaussian prior is assumed on the Fourier coefficients  $\mathbf{a}$ , and uniform priors assumed on the red noise amplitude hyper parameters, those amplitude parameters can be marginalised over analytically in exactly the same manner as described in Section 5.3.

### 5.5 Including dispersion measure variations

The plasma located in the interstellar medium (ISM), as well as in solar winds and the ionosphere can result in delays in the propagation of the pulse signal between the pulsar and the observatory, an effect that appears as a red noise signal in the timing residuals.

The severity of the observed dispersion measure variations, however, is dependent upon the observing frequency, and as such we can use this additional information to isolate this component of the red noise from effects that do not have this dependence, such as red spin noise.

In particular, the group delay  $t_g(\nu)$  for an observing frequency  $\nu$  is given by the relation:

$$t_g(\nu) = K \text{ DM} / (\nu^2), \quad (50)$$

where the dispersion constant  $K$  is given by:

$$K \equiv 4.15 \times 10^{15} \text{ Hz}^2 \text{ cm}^3 \text{ pc}^{-1} \text{ s} \quad (51)$$

and the dispersion measure is defined as the integral of the electron density  $n_e$  from the Earth to the pulsar:

$$\text{DM} = \int_0^L n_e dl. \quad (52)$$

Dispersion measure corrections can be included in the analysis as an additional set of stochastic parameters in almost the same way as the frequency independent spin noise. We begin by first defining a vector  $\mathbf{D}$  of length equal to the number of pulse profiles for a given pulsar as:

$$D_i = K/(\nu_i^2) \quad (53)$$

for observation  $i$  with observing frequency  $\nu_i$ .

We then write the basis vectors that describe the dispersion measure Fourier modes as:

$$F_{DM}(\nu, t_i) = \frac{1}{T} \sin(2\pi\nu t_i) D_i \quad (54)$$

and an equivalent cosine term, where  $T$  is the length of the observing timespan, and  $\nu$  denotes the frequency of the signal to be parameterised as before, where the set of frequencies to be included is defined in the same way as for the red spin noise. Unlike when modelling the red spin noise, we no longer have the quadratic in the timing model to act as a proxy to the low frequency ( $\nu < 1/T$ ) DM variations in our data. As such these terms must be accounted for either by explicitly including these low frequencies in the model, or by including a quadratic in DM to act as a proxy, as with the red noise, defined as:

$$Q_{DM}(t_i) = \delta_0 t_i D_i + \delta_1 t_i^2 D_i, \quad (55)$$

with  $\delta_{0,1}$  free parameters to be fit for, and  $t_i$  the barycentric arrival time for TOA  $i$ . This is most simply done by including these terms in the set of timing model parameters  $\epsilon$ . If these terms are not included in the model then power from frequencies lower than  $1/T$  will be absorbed by the Fourier coefficients included in  $\mathbf{F}_{DM}$ , biasing the estimated power spectrum.

As with the red noise we can then include the DM signal realisation in our model TOAs:

$$\hat{\tau}(\epsilon, \mathbf{a}_{\text{red}}, \mathbf{a}_{DM}) = \tau(\epsilon) - \mathbf{F}_{\text{red}} \mathbf{a}_{\text{red}} - \mathbf{F}_{DM} \mathbf{a}_{DM}, \quad (56)$$

where we have factored the quadratic  $Q_{DM}$  into the timing model  $\tau(\epsilon)$ . Finally we then define the matrix of DM power spectrum coefficients  $\boldsymbol{\varphi}_{DM}$  such that:

$$\text{Pr}(\mathbf{a}_{DM} | \boldsymbol{\rho}_{DM}) \propto \frac{1}{\sqrt{\det \boldsymbol{\varphi}_{DM}}} \exp \left[ -\frac{1}{2} \mathbf{a}_{DM}^* \boldsymbol{\varphi}_{DM}^{-1} \mathbf{a}_{DM} \right]. \quad (57)$$

We note here that, as with the red noise, if desired additional terms can be added into the DM Fourier matrix  $\mathbf{F}_{DM}$  to model, for example, additional annual variations in the data, and that as before the DM spectrum coefficients can be parameterised with a non-Gaussian prior.

## 5.6 A non-Gaussian stochastic gravitational wave background

The final application to pulsar timing of the non-Gaussian formalism developed thus far that we will consider is to a stochastic gravitational wave background (GWB) for which we follow the approach given in (Lentati et al. 2013).

As with the intrinsic red noise model described in section 5.4 we parameterise the signal in each pulsar using the Fourier basis given in equation 42. When dealing with a signal from a GWB, however, it is crucial to include the cross correlated signal between the pulsars on the sky. We do this by using the Hellings-Downs relation Hellings & Downs (1983):

$$\begin{aligned} \Theta_{mn} &= \frac{3}{2} \frac{1 - \cos(\theta_{mn})}{2} \ln \left( \frac{1 - \cos(\theta_{mn})}{2} \right) \\ &- \frac{1}{4} \frac{1 - \cos(\theta_{mn})}{2} + \frac{1}{2} + \frac{1}{2} \delta_{mn}, \end{aligned} \quad (58)$$

where  $\theta_{mn}$  is the angle between the pulsars  $m$  and  $n$  on the sky and  $\Theta_{mn}$  represents the expected correlation between the TOAs given

an isotropic background. With this addition our covariance matrix for the Fourier coefficients becomes

$$\varphi_{mi,nj} = \langle a_{mi} a_{nj}^* \rangle = \Theta_{mn} \varphi_i \delta_{ij}, \quad (59)$$

where there is no sum over  $i$ , which results in a band diagonal matrix for which calculating the inverse is extremely computationally efficient.

Our prior term  $\text{Pr}(\mathbf{a}|\boldsymbol{\varphi})$  is then similar to that given in Eq. 47, with the only difference that  $\mathbf{a}$  is now the concatenated vector of Fourier coefficients for all pulsars, and  $\boldsymbol{\varphi}$  includes the correlations between pulsars. We can then parameterise the non-Gaussianity at each GWB frequency as in Eq. 49, however with multiple pulsars we need not assume a power law prior on the coefficients if desired.

## 6 APPLICATION TO SIMULATIONS

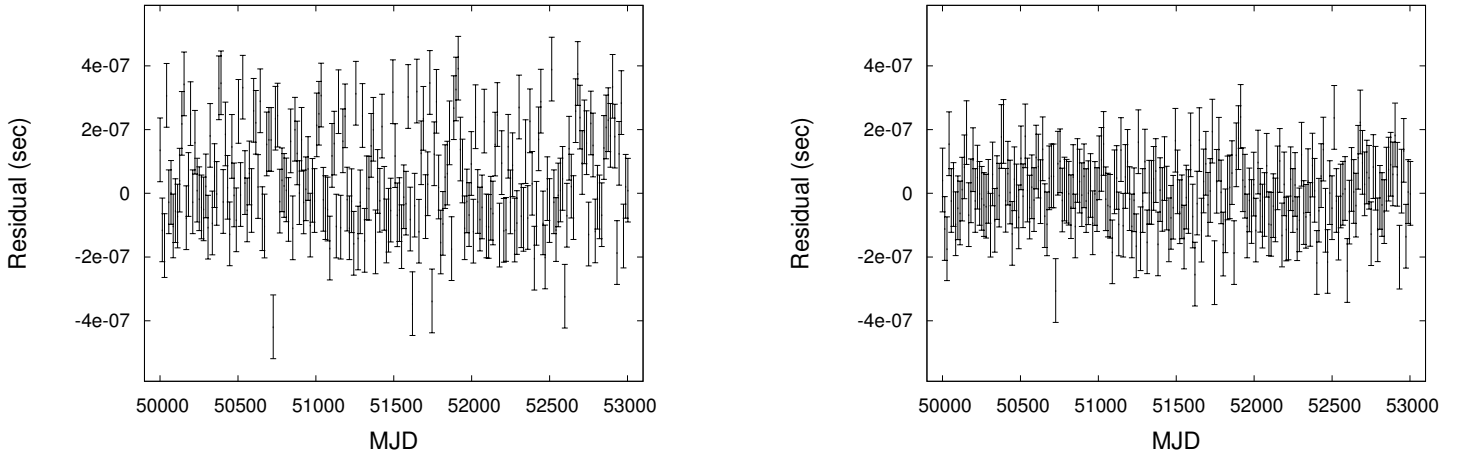
We now apply the likelihood developed in section 5 to two simulations. Simulation 1 includes non-Gaussian noise with the same distribution as in the toy model in section 4, and Simulation 2 includes only Gaussian noise. In both cases we set the width of the ground state Gaussian distribution to be  $10^{-6}$  seconds. We then simulate an  $\sim 8$  year dataset for the isolated pulsar J0030+0451 with observations spaced  $\sim 2$  weeks apart resulting in a total of 216 TOAs onto which we add a noise realisation drawn from the two distributions using the same seed both times. The timing model parameters used in the simulations are listed in Table 2, and are based on values provided as part of the first IPTA data challenge. The timing residuals that remain from the two simulations after subtracting the injected timing model are shown in Fig. 2.

In both simulations we will compare the parameter estimates obtained using two models. Model 1 will include the parameters  $\alpha_{1,3}$  to model any deviations from Gaussianity that the timing residuals might be subject to, and in Model 2 will set these parameters equal to zero, therefore constraining the probability density to be purely Gaussian.

While in principle the  $\alpha$  coefficients can be real or complex valued, as they only appear in the likelihood in the term  $|\sum_n^{\alpha_{\text{max}}} \alpha_n C_n H_n|^2$  we restrict our search to only real values in order to eliminate potential degeneracies that would otherwise arise. In both simulations we take the prior range on these coefficients to be  $[-1, 1]$ .

### 6.1 Simulation 1

Table 2 lists the mean parameter estimates and standard deviations for the two models applied to simulation 1. In addition, Fig. 3 shows the one-dimensional marginalised posteriors for the timing model parameters and the white noise scaling parameter  $\beta$  for model 1 (red solid lines) and model 2 (blue dotted lines) respectively. In all cases the constraints on the timing model parameters are improved by a factor  $\sim 2$  when we include the additional parameters and thus correctly account for the non-Gaussian nature of the noise. The degree to which the timing model parameters are affected will, however, naturally depend on the severity of the non-Gaussianity in the residuals. In addition from Fig. 2 we see that the effect of the non-Gaussian terms is to i) increase the number of outliers, and ii) to shift the mean of the residuals as the distribution is no longer symmetric. The result of this is that when assuming a Gaussian likelihood the constant phase term is significantly offset from zero, and the EFAC term  $\beta$  is increased in order to accommodate the outliers. While the shift in the phase term is irrelevant, the increase in



**Figure 2.** Non-Gaussian (left) and Gaussian (right) timing residuals from simulations one and two respectively after subtracting the simulated timing model.

the EFAC term leads directly to the decrease in sensitivity to the timing model parameters observed previously. Comparing the evidence between the two models we find that the additional parameters in model 1 are strongly favoured with  $\Delta \log E = 18$ , suggesting a definitive detection of non-Gaussianity in the residuals.

## 6.2 Simulation 2

As for simulation 1 we list the mean parameter estimates and standard deviations using models 1 and 2 for simulation 2 in Table 2. Similarly Fig. 4 shows the one-dimensional marginalised posteriors for the timing model parameters and the white noise scaling parameter  $\beta$  for model 1 (red solid lines) and model 2 (blue dotted lines) respectively. In this simulation, with the exception of the phase offset, the posteriors for the timing model parameters are identical for the two models, as could be expected as the noise is now Gaussian in nature. The phase offset however is now totally unconstrained across the prior in model 1, and the uncertainties in the scaling parameter  $\beta$  are a factor  $\sim 4$  greater than in model 2.

We can understand this disparity by looking at the two-dimensional marginalised posteriors for the phase offset,  $\beta$ , and  $\alpha$  parameters in Fig. 5. Here it becomes clear that the phase offset and  $\alpha_1$  parameter are in this instance completely correlated. Qualitatively this simply represents that there is no difference between a dataset with a phase offset, and a dataset whose noise probability density is both symmetric and offset from zero. We also see that there is a strong correlation between the  $\beta$  scaling parameter and  $\alpha_2$  leading to the increased uncertainties in the former relative to model 2.

Comparing the evidence between the two models we find that, as expected, model 2 is favoured with  $\Delta \log E = 4$  providing strong support for the simpler model.

## 7 APPLICATION TO REAL DATA

We now apply the likelihood developed in section 5 to the publicly available Parkes Pulsar Timing Array (PPTA) data release 1 dataset for the binary pulsar J0437–4715 (Manchester et al. 2013). In brief, this dataset spans 15 years of observations, with approximately the

first nine consisting of only single frequency, 20cm observations, and the remainder including additional 10cm and 50cm observations. The timing model we fit includes the 21 parameters listed in Table 4 in addition to a set of 12 “jumps” (offsets between different observing systems). We then simultaneously fit for a power law red noise process and power law dispersion measure variations, including periods from the length of the dataset  $T$ , down to one month, at intervals of  $1/T$  as described in sections 5.4 and 5.5 respectively. Finally we also include an EFAC and EQUAD parameter for each observing system group. In order to explore the effects of potential non-Gaussianity in the dataset we then consider two cases. Firstly, that both the radiometer noise term ( $\beta\sigma$ ) and quadrature term ( $\gamma$ ) describe Gaussian random processes, and secondly, that while the radiometer noise is Gaussian, the quadrature term is non-Gaussian. In the latter case we parameterise this non-Gaussianity as in Section 5.1 where we fit for the convolved PDF of the two distributions, rather than sampling from the  $j$  parameters numerically.

Residuals for this dataset after subtracting the maximum likelihood Gaussian timing solution given in Table 4 are shown in Fig. 6 (top), and after also subtracting the maximum likelihood red noise and dispersion measure variations (middle). In the bottom left panel we show the normalised residuals after dividing each point by its error bar, and in the bottom right panel we show a histogram of these normalised residuals, overlaid with a unit Gaussian to show the expected number counts in each bin.

From the plot of the normalised residuals it is clear that there are a significant number of outliers ( $> 4 - 5\sigma$ ) compared to the number expected from a Gaussian distribution. In addition the histogram suggests that there is also an over abundance of points at small deviations ( $< 0.5\sigma$ ). As seen in Section 6 this behaviour is indicative of the white noise parameters in the Gaussian fit overestimating the error bars of the points in order to best accommodate outliers.

### 7.1 A post fit evidence comparison

As in section 6 we would like to compute the difference in the Evidence for the Gaussian and non-Gaussian models, however for this dataset the total dimensionality of the problem is  $\sim 800$ , which makes the calculation of the evidence both extremely expensive



**Table 2.** Parameter estimates for the two simulations of PSR J0030+0451. Figures in parentheses represent one standard deviation in the least-significant digits quoted.

Simulation 1			
Model Parameter	Simulation	Model 1	Model 2
log Evidence	-	3052.9	3035.0
Right ascension, $\alpha$	00:30:27.4299630	00:30:27.429956(7)	00:30:27.429964(14)
Declination, $\delta$	+04:51:39.75230	+04:51:39.7525(3)	+04:51:39.7522(5)
Pulse frequency, $\nu$ ( $\text{s}^{-1}$ )	205.53069608827310	205.53069608827315(6)	205.53069608827300(12)
First derivative of pulse frequency, $\dot{\nu}$ ( $\text{s}^{-2}$ )	-1.3061388e-16	-1.306139(5) $\times 10^{-16}$	-1.306149e-16(9) $\times 10^{-16}$
Proper motion in right ascension, $\mu_\alpha$ ( $\text{mas yr}^{-1}$ )	-4.054	-4.09(2)	-4.05(4)
Proper motion in declination, $\mu_\delta$ ( $\text{mas yr}^{-1}$ )	-5.03	-4.94(5)	-5.0(1)
Parallax, $\pi$ (mas)	4.023	4.037(14)	4.02(3)
$\beta$	1	1.00(4)	1.62(8)
$\alpha_1$	0.1	0.13(5)	-
$\alpha_2$	0.1	0.15(5)	-
$\alpha_3$	0.4	0.41(4)	-

Simulation 2			
Model Parameter	Simulation	Model 1	Model 2
log Evidence	-	3131.5	3136.4
Right ascension, $\alpha$	00:30:27.4299630	00:30:27.429964(9)	00:30:27.429964(9)
Declination, $\delta$	+04:51:39.75230	+04:51:39.7523(3)	+04:51:39.7523(3)
Pulse frequency, $\nu$ ( $\text{s}^{-1}$ )	205.53069608827310	205.53069608827304(7)	205.53069608827304(7)
First derivative of pulse frequency, $\dot{\nu}$ ( $\text{s}^{-2}$ )	-1.3061388e-16	-1.306146e-16(6) $\times 10^{-16}$	-1.306146e-16(6) $\times 10^{-16}$
Proper motion in right ascension, $\mu_\alpha$ ( $\text{mas yr}^{-1}$ )	-4.054	-4.06(3)	-4.06(3)
Proper motion in declination, $\mu_\delta$ ( $\text{mas yr}^{-1}$ )	-5.03	-5.01(6)	-5.01(6)
Parallax, $\pi$ (mas)	4.023	4.037(14)	4.025(16)
$\beta$	1	1.1(2)	0.99(5)
$\alpha_1$	0	-0.0(3)	-
$\alpha_2$	0	-0.02(12)	-
$\alpha_3$	0	-0.01(7)	-

computationally, and also much less precise numerically than in the lower dimensional simulation.

As such we consider that even in the presence of significant non-Gaussianity, the timing model parameter estimates obtained from the Gaussian analysis did not differ significantly from the non-Gaussian analysis, only their uncertainties changed. As such, in order to obtain an approximate value for the evidence we can use the post-fit residuals, after subtracting the maximum likelihood timing model, red noise and dispersion measure variations and then fit only for the EFAC and EQUAD parameters in those residuals and an additional offset term. This decreases the dimensionality to  $\sim 30$ , allowing us to use MULTINEST to compare the evidence for the Gaussian and non-Gaussian cases as before. When including non-Gaussian coefficients in our model, both in this test and in the subsequent full analysis, we exclude the  $\alpha_1$  term in order to minimise the covariance between the offset and the non-Gaussian parameters.

Table 3 lists the log evidence values for different sets of non-Gaussian coefficients, normalised such that the log evidence for no additional coefficients (i.e. assuming Gaussian statistics) is 0. We see that there is a significant increase in the log evidence ( $\sim 39$ ) when including even just two coefficients, indicating definitive support for their inclusion in the model. As the number increases the rise in evidence increases, reaching a maximum with 4 included coefficients. Given the timing model, red noise and dispersion measure variation solutions that were subtracted from the data were

**Table 3.** Log evidence values for different numbers of non-Gaussian coefficients in a post fit analysis of the J0437–4715 residuals.


non-Gaussian Coefficients Included	log Evidence
0	0
$\alpha_{2..3}$	38.7
$\alpha_{2..4}$	47.4
$\alpha_{2..5}$	48.3
$\alpha_{2..6}$	48.2

obtained from a Gaussian analysis, we will however still include coefficients up to and including  $\alpha_6$  in the full analysis.

## 7.2 Joint Bayesian analysis

Given the large dimensionality of the problem this analysis cannot be carried out using MULTINEST. As such we make use of the 'Guided Hamiltonian Sampler' used previously in pulsar timing analysis in (Lentati et al. 2013). This sampler makes use of both gradient information in the likelihood, and also the hessian in order to efficiently sample from large parameter spaces.

Table 4 lists the timing model parameter estimates and their nominal standard deviations for both the Gaussian and non-Gaussian analysis. In all cases we find the parameter estimates and their uncertainties to be consistent between both methods. In Fig. 7 (top) we show the one and two-dimensional marginalised poste-




**Figure 3.** 1-dimensional marginalised posteriors for the timing model and  $\beta$  parameters in simulated dataset 1 for the isolated pulsar PSR J0030+0451 for model 1 (red solid line) and model 2 (blue dotted line). Values on the  $x$ -axes for the timing model parameters are given in terms of the standard deviation in that parameter returned by the analysis when including the additional terms, with the injected parameter value at 0 in all cases. The non-Gaussian nature of the noise results in a significant increase in the scaling parameter  $\beta$  in model 2, leading to a decrease in the precision with which the timing model parameters are detected by a factor  $\sim 2$  relative to model 1.

rrior distributions for the red noise and dispersion measure variation power law amplitudes and spectral indices for the non-Gaussian (left) and Gaussian (right) analysis. Both are also extremely consistent with one another, however when overlaying the two sets of 1-dimensional posterior distributions for each of the 4 parameters separately (bottom 4 panels) some differences become apparent between the non-Gaussian (blue dashed lines) and Gaussian (red solid lines) analysis. In particular the dispersion measure variation power law parameter estimates show a slight shift towards higher amplitudes and shallower spectral indices in the non-Gaussian case.

Despite these similarities in the timing and stochastic parameter estimates between the Gaussian and non-Gaussian analysis, Fig. 8 indicates a definitive detection of non-Gaussianity in the dataset, in agreement with the difference in the log evidence for the noise only analysis. In the top plot we show the one and two-dimensional marginalised posterior distributions for the 5 non-Gaussian coefficients fit in the analysis of J0437–4715. Vertical lines are included at 0 where visible in the plots, however, except for  $\alpha_5$  all the coefficients are inconsistent with this value. In the bottom plot we

then show the set of equally weighted PDFs obtained from the non-Gaussian analysis (black lines) setting  $\gamma = 1$ . In addition we over plot the mean of the distribution (red line) and a unit Gaussian (blue line) all of which have been normalised to have a sum of 1. The difference between the Gaussian and non-Gaussian PDFs is clear, with a larger probability for both small ( $|\sigma| < 1$ ) and larger ( $|\sigma| > 4$ ) deviations than given by the Gaussian PDF. That such a significant detection of non-Gaussianity does not lead to larger changes in the parameter estimates can potentially be attributed to a frequency dependence on the significance of the  $\alpha$  parameters. In Shannon et al. (2014) the 10cm J0437–4715 data was found to be describable through Gaussian statistics alone. This would suggest that the non-Gaussianity we detect exists primarily at low frequencies. In Figure 9 we show the normalised residuals from Fig. 7 separated into its 10cm, 20cm and 50cm components, along with histograms for each wavelength. Here the increase in non-Gaussian behaviour can clearly be seen as the wavelength increases. Given the lowest frequencies have the greatest degree of non-Gaussianity it is less surprising that there is little impact on the timing or red spin noise



**Figure 4.** 1-dimensional marginalised posteriors for the timing model and  $\beta$  parameters in simulated dataset 2 for the isolated pulsar PSR J0030+0451 for model 1 (red solid line) and model 2 (blue dotted line). Values on the  $x$ -axes for the timing model parameters are given in terms of the standard deviation in that parameter returned by the analysis when including the additional terms, with the injected parameter value at 0 in all cases. Here the simulated noise was Gaussian in nature, and so with the exception of the phase offset, the parameter estimates for the timing model parameters using the two models are identical. The disparity in the phase offset and  $\beta$  parameters is discussed in Section 6.2.

parameters, as the low frequency data contributes the least to these parts of the model. The low frequencies do, however, contribute greatly to the constraints on dispersion measure variations, and it is here we see the greatest difference between the Gaussian and non-Gaussian models.

## 8 CONCLUSION

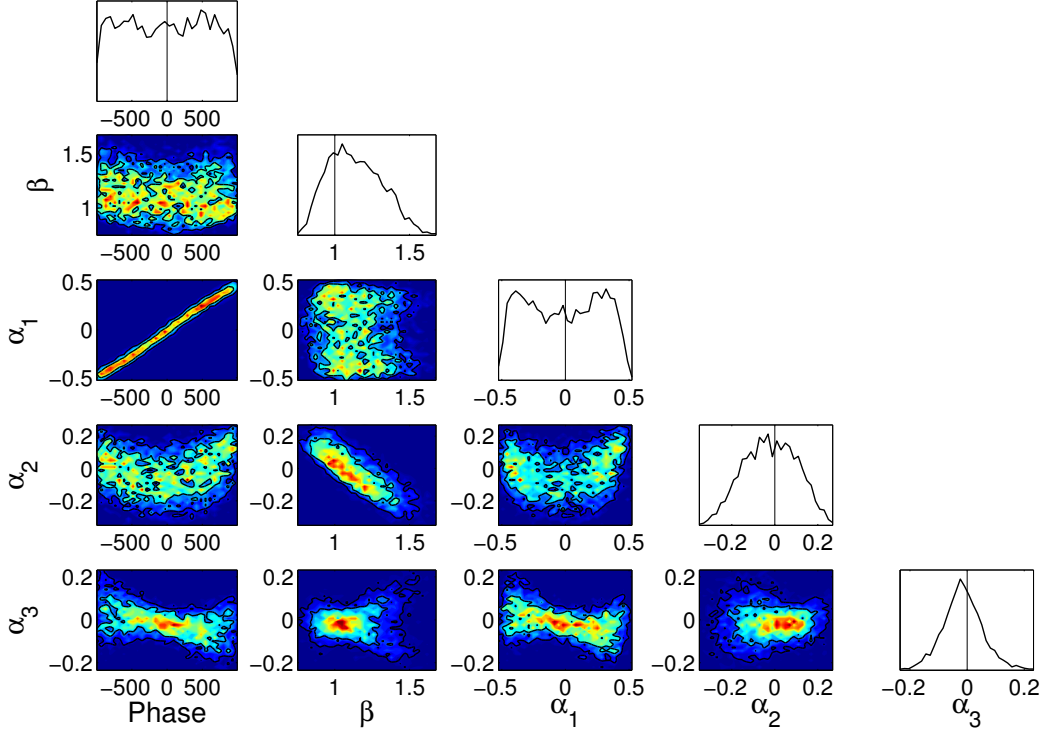
In this paper we have introduced a method of performing a robust Bayesian analysis of non-Gaussianity present in the residuals in pulsar timing analysis, simultaneously with the pulsar timing model, and additional stochastic parameters such as those describing the red noise, and dispersion measure variations present in the data. Deviations from Gaussianity are described using a set of parameters  $\alpha$  that act to modify the probability density of the noise, such that  $\alpha = 0$  describes Gaussian noise, and any non zero values provide support for non-Gaussian behaviour. The advantage of this method is that one may use a finite set of non-zero  $\alpha_m$  to model the non-Gaussianity, without mathematical inconsistency. Any trunca-

tion of the series still yields a proper distribution, in contrast to the more commonly used Edgeworth expansion (e.g. Contaldi et al. (2000)).

We applied this method to two simulated datasets. In simulation one the noise was drawn from a non-Gaussian distribution, and in simulation 2 it was purely Gaussian. In simulation 1, the effect of the non-Gaussianity was to introduce a higher proportion of outliers relative to a Gaussian distribution. This resulted in an overestimation of the TOA uncertainties when assuming a Gaussian likelihood, and decreased the precision with which the timing model parameters could be extracted compared to an analysis that correctly incorporated the non-Gaussian behaviour on the noise.

In the second case we showed that the parameter estimates of the timing model parameters of interest were consistent when including, or not, the  $\alpha$  parameters, as is to be expected when the noise is Gaussian.

We then applied this method to the publicly available Parkes Pulsar Timing Array (PPTA) data release 1 dataset for the binary pulsar J0437–4715. We detect a significant non-Gaussian compo-

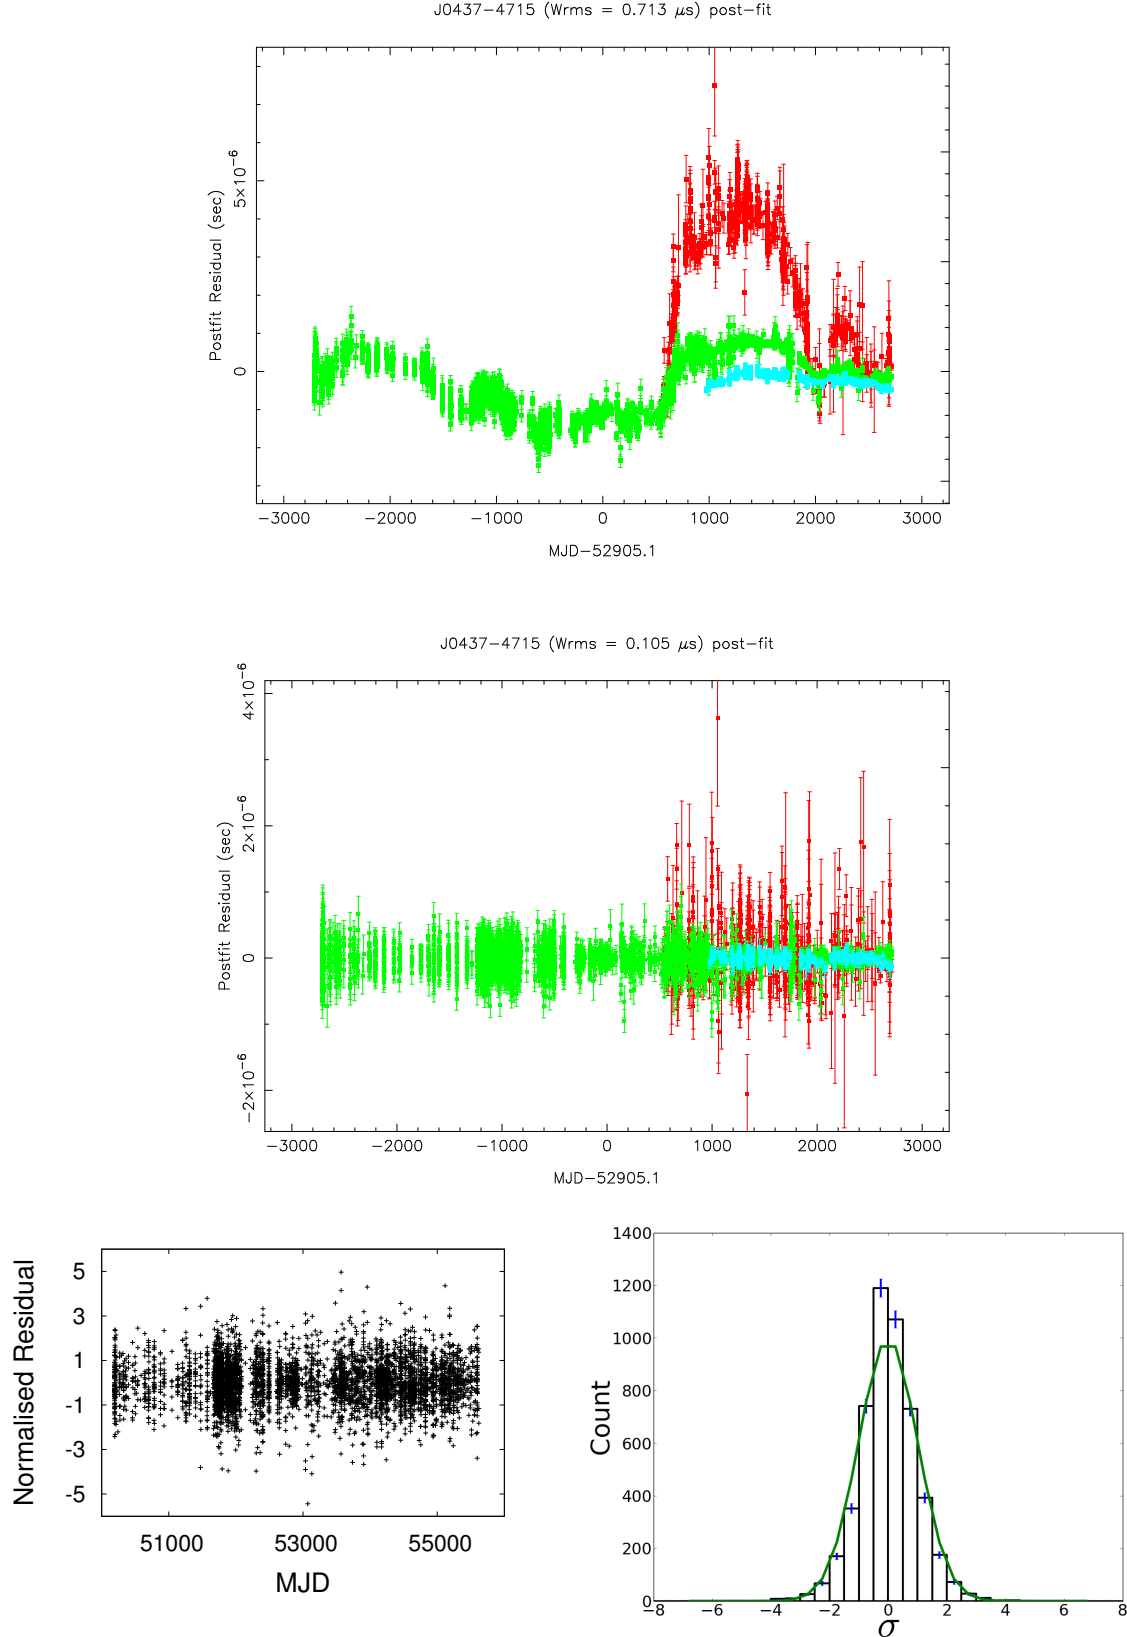


**Figure 5.** One and two-dimensional marginalised posterior distributions for the phase offset,  $\beta$ , and  $\alpha$  parameters in simulation 2. The phase offset and  $\alpha_1$  parameters are in this instance completely correlated. Qualitatively this simply represents that there is no difference between a dataset with a phase offset, and a dataset whose noise probability density is both symmetric and offset from zero. Similarly  $\beta$  and  $\alpha_2$  show strong correlations, leading to the increased uncertainty in the parameter seen in Fig. 4.


nent in the non-thermal component of the uncorrelated noise, however as the non-Gaussianity is most dominant in the lowest frequency data the impact on the timing precision in the pulsar is minimal, with only the parameter estimates of the power law dispersion measure variations being visible changed between the Gaussian and non-Gaussian analysis.

## REFERENCES

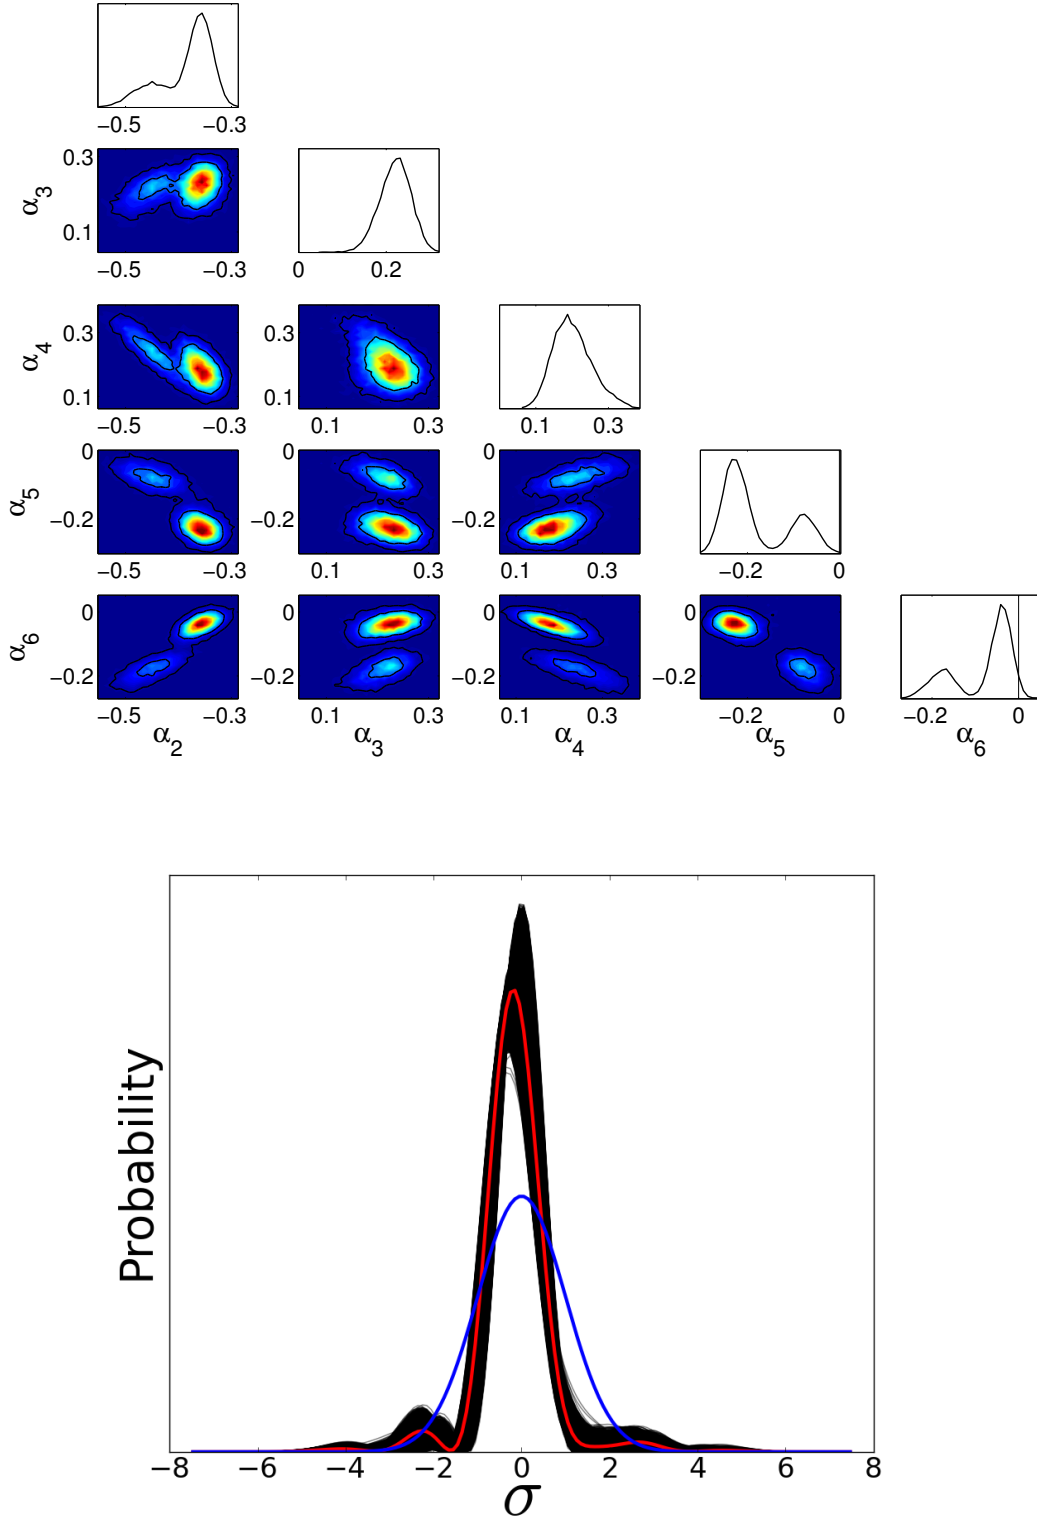
- Coles W., Hobbs G., Champion D. J., Manchester R. N., Verbiest J. P. W., 2011, *MNRAS*, 418, 561
- Contaldi C. R., Ferreira P. G., Magueijo J., Górski K. M., 2000, *ApJ*, 534, 25
- Demorest P. B., et al., 2013, *ApJ*, 762, 94
- Edwards R. T., Hobbs G. B., Manchester R. N., 2006, *MNRAS*, 372, 1549
- Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
- Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
- Hall P., 1989, *Ann. Statist.* Volume 17, no. 2, 589–605
- Hellings R. W., Downs G. S., 1983, *ApJ*, 265, L39
- Hobbs G., et al., 2009, *MNRAS*, 394, 1945
- Hobbs G. B., Edwards R. T., Manchester R. N., 2006, *MNRAS*, 369, 655
- Jaffe A. H., Backer D. C., 2003, *ApJ*, 583, 616
- Janssen G. H., Stappers B. W., Kramer M., Purver M., Jessner A., Cognard I., 2008, in Bassa C., Wang Z., Cumming A., Kaspi V. M., eds, *AIP Conf. Proc.* Vol. 983, 40 Years of Pulsars: Millisecond Pulsars, Magnetars and More. Am. Inst. Phys., New York, p. 633
- Kaspi V. M., Taylor J. H., Ryba M. F., 1994, *ApJ*, 428, 713
- Kramer M., et al., 2006, *Science*, 314, 97
- Lee K. J., Bassa C. G., Janssen G. H., Karuppusamy R., Kramer M., Smits R., Stappers B. W., 2012, *MNRAS*, 423, 2642
- Lentati L., Alexander P., Hobson M. P., Feroz F., van Haasteren R., Lee K. J., Shannon R. M., 2014, *MNRAS*, 437, 3004
- Lentati L., Alexander P., Hobson M. P., Taylor S., Gair J., Balan S. T., van Haasteren R., 2013, *Phys. Rev. D*, 87, 104021
- Lorimer D. R., Kramer M., 2004, Ellis R., Huchra J., Kahn S., Rieke G., Stetson P. B., eds, *Handbook of Pulsar Astronomy*. Cambridge Univ. Press, Cambridge
- Manchester R. N., et al., 2013, *PASA*, 30, 17
- Matsakis D. N., Taylor J. H., Eubanks T. M., 1997, *A&A*, 326, 924
- O’Ruanaidh J. J. K., Fitzgerald W. J., 1996, *Numerical Bayesian Methods Applied to Signal Processing*. Springer-Verlag, New York
- Phinney E. S., 2001, preprint (astro-ph/0108028)
- Rocha G., Magueijo J., Hobson M., Lasenby A., 2001, *Phys. Rev. D*, 64, 063512
- Shannon R. M., Cordes J. M., 2010, *ApJ*, 725, 1607
- Shannon R. M., et al., 2014, *MNRAS*, 443, 1463
- Skilling J., 2004, in Fischer R., Preuss R., von Toussaint U., eds, *AIP Conf. Proc.* Vol. 735, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. Am. Inst. Phys., New York, p. 395
- Taylor J. H., Weisberg J. M., 1989, *ApJ*, 345, 434



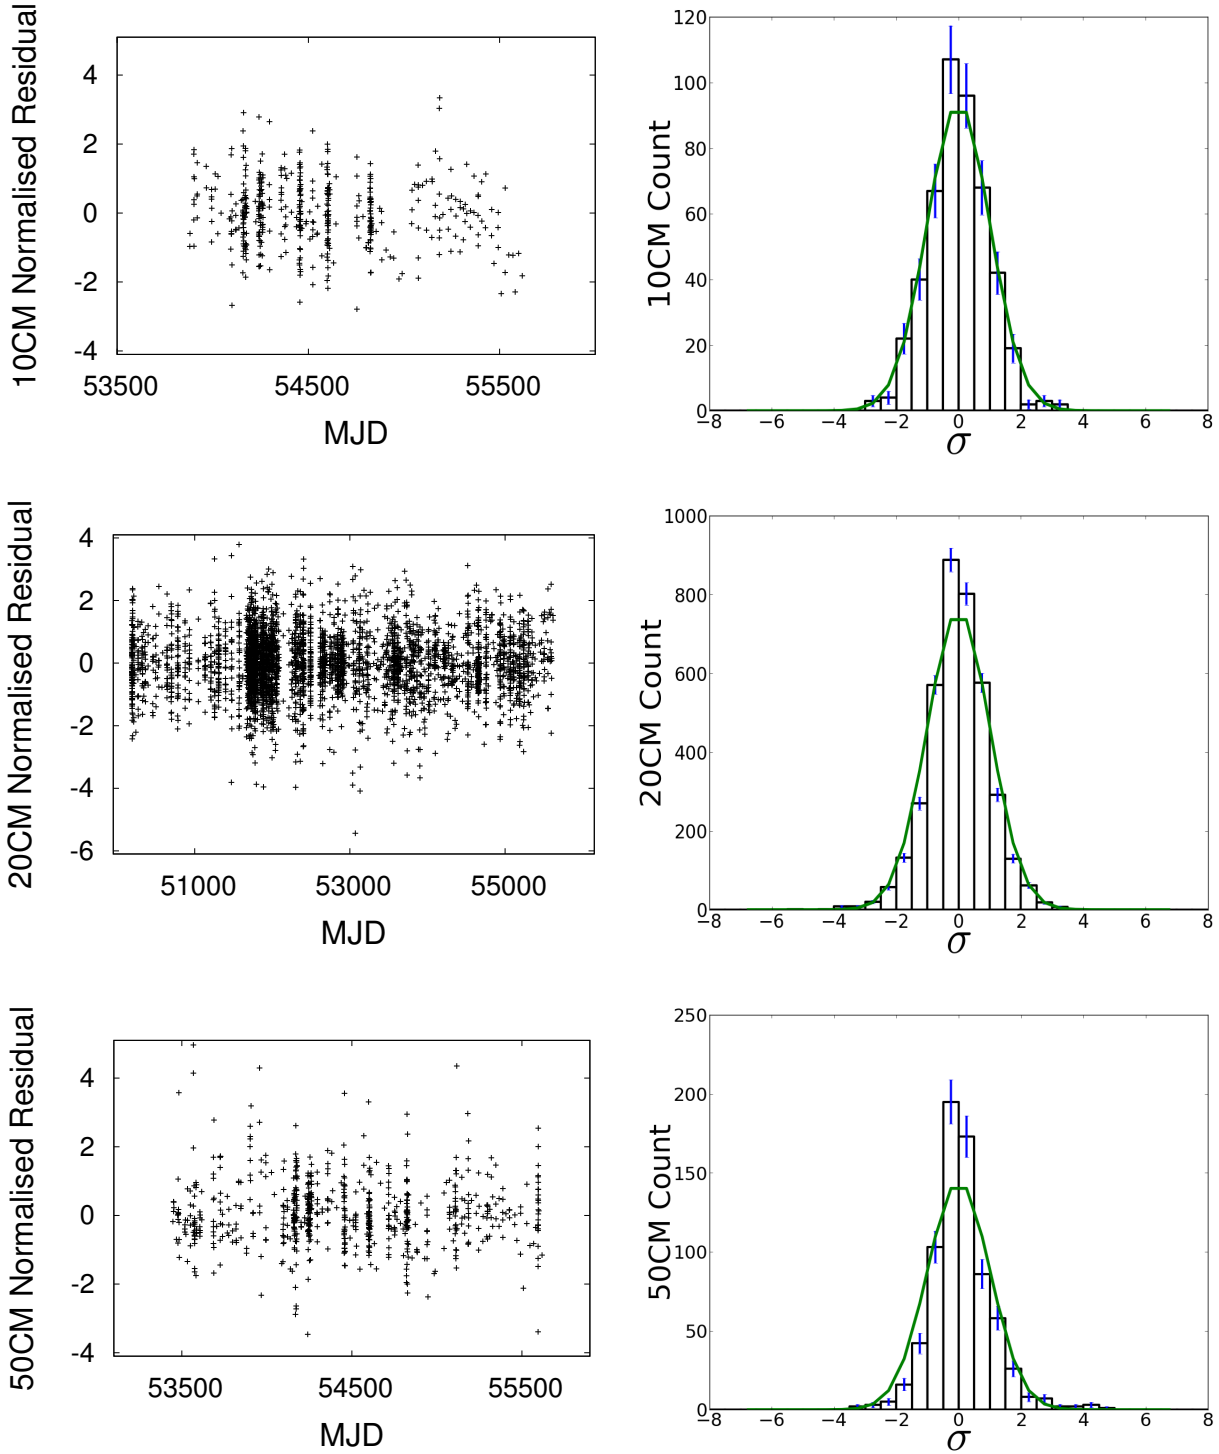
**Figure 6.** Residuals for the publicly available PPTA data release 1 dataset for the binary pulsar J0437-4715 after subtracting the maximum likelihood Gaussian timing solution given in Table 4 (top), and additionally after subtracting the maximum likelihood red noise and dispersion measure variations (middle). Colours indicate 10cm (blue), 20cm (green) and 50cm (red) observing wavelengths. In the bottom left panel we show the normalised residuals after dividing each point by its error bar, and in the bottom right panel we show a histogram of these normalised residuals, overlaid with a unit Gaussian to show the expected number counts in each bin. Error bars for each histogram bin are given by  $\sqrt{N}$  with  $N$  the number of points in the bin.



**Figure 7.** (Top) One and two-dimensional marginalised posterior distributions for the red noise and dispersion measure variation power law amplitudes and spectral indices for the non-Gaussian (left) and Gaussian (right) analysis. Both are extremely consistent with one another, however when overlaying the two sets of 1-dimensional posterior distributions for each of the 4 parameters separately (bottom 4 panels) some differences become apparent between the non-Gaussian (blue dashed lines) and Gaussian (red solid lines) analysis. In particular the dispersion measure variation power law parameter estimates show a slight shift towards higher amplitudes and shallower spectral indices in the non-Gaussian case.



**Figure 8.** (Top) One and two-dimensional marginalised posterior distributions for the 5 non-Gaussian coefficients fit in the analysis of J0437–4715. Vertical lines are included at 0 where visible in the plot, however, except for  $\alpha_5$  all the coefficients are inconsistent with this value suggesting a definitive detection of non-Gaussianity in the dataset. (Bottom) The set of equally weighted probability density functions obtained from the non-Gaussian analysis of J0437–4715 (black lines) with  $\gamma = 1$ . In addition we show the mean of the distribution (red line) and a unit Gaussian (blue line) all normalised to have a sum of 1. The difference between the Gaussian and fitted non-Gaussian functions is clear, with a larger probability for both small ( $|\sigma| < 1$ ) and larger ( $|\sigma| > 4$ ) deviations than given by the Gaussian probability function.



**Figure 9.** (Left) Normalised Residuals for the publicly available PPTA data release 1 dataset for the binary pulsar J0437–4715 after subtracting the maximum likelihood Gaussian timing solution given in Table 4 and the maximum likelihood red noise and dispersion measure variations for 10CM (top), 20CM (middle), and 50CM (bottom) datapoints. (Right) Histogram of the normalised residuals for 10CM (top), 20CM (middle), and 50CM (bottom) datapoints, overlaid with a unit Gaussian to show the expected number counts in each bin. Error bars for each histogram bin are given by  $\sqrt{N}$  with  $N$  the number of points in the bin.



**Table 4.** Parameters for PSR J0437–4715. Figures in parentheses are the nominal standard deviations in the least-significant digits quoted.

Fit and data-set		
Pulsar name . . . . .	J0437–4715	
MJD range . . . . .	50191.0–55619.2	
Data span (yr) . . . . .	14.86	
Number of TOAs . . . . .	5052	
Measured Quantities		
Model Parameter	Non–Gaussian	Gaussian
Right ascension, $\alpha$ (rad) . . . . .	1.20979650940(10)	1.20979650943(11)
Declination, $\delta$ (rad) . . . . .	–0.82471224153(8)	–0.82471224154(8)
Pulse frequency, $\nu$ ( $\text{s}^{-1}$ ) . . . . .	173.6879458121850(3)	173.6879458121849(4)
First derivative of pulse frequency, $\dot{\nu}$ ( $\text{s}^{-2}$ ) . . . . .	–1.728365(4) $\times 10^{-15}$	–1.728365(4) $\times 10^{-15}$
Dispersion measure, DM ( $\text{cm}^{-3}\text{pc}$ ) . . . . .	2.64462(11)	2.64461(11)
First derivative of dispersion measure, $\dot{DM}$ ( $\text{cm}^{-3}\text{pc yr}^{-1}$ ) . . . . .	–6(6) $\times 10^{-5}$	–7(7) $\times 10^{-5}$
DM2 ( $\text{cm}^{-3}\text{pc yr}^{-2}$ ) . . . . .	–1(2) $\times 10^{-6}$	–1(2) $\times 10^{-6}$
Proper motion in right ascension, $\mu_\alpha \cos \delta$ ( $\text{mas yr}^{-1}$ ) . . . . .	121.439(3)	121.441(3)
Proper motion in declination, $\mu_\delta$ ( $\text{mas yr}^{-1}$ ) . . . . .	–71.474(3)	–71.474(3)
Parallax, $\pi$ (mas) . . . . .	6.4(2)	6.3(2)
Orbital period, $P_b$ (d) . . . . .	5.7410462(3)	5.7410461(3)
Epoch of periastron, $T_0$ (MJD) . . . . .	54530.1722(3)	54530.1721(3)
Projected semi-major axis of orbit, $x$ (lt-s) . . . . .	3.36671463(8)	3.36671464(8)
Longitude of periastron, $\omega_0$ (deg) . . . . .	1.35(2)	1.36(2)
Orbital eccentricity, $e$ . . . . .	1.91800(14) $\times 10^{-5}$	1.91796(15) $\times 10^{-5}$
First derivative of orbital period, $\dot{P}_b$ . . . . .	3.724(6) $\times 10^{-12}$	3.724(6) $\times 10^{-12}$
First derivative of $x$ , $\dot{x}$ ( $10^{-12}$ ) . . . . .	1(2) $\times 10^{-15}$	1(2) $\times 10^{-15}$
Periastron advance, $\dot{\omega}$ (deg/yr) . . . . .	0.0150(12)	0.0150(13)
Companion mass, $M_c$ ( $M_\odot$ ) . . . . .	0.223(14)	0.223(15)
Longitude of ascending node, $\Omega$ (degrees) . . . . .	208.0(12)	208.3(13)
Orbital inclination angle, $i$ (degrees) . . . . .	137.1(8)	137.3(8)

van Haasteren R., et al., 2011, MNRAS, 414, 3117